

Critical Capabilities for Insight Engines

Updated 4 February 2022, Published 30 March 2021 - ID G00465453 - 35 min read

UPDATED This research has been updated to include recent changes in capability.

By Stephen Emmott, Anthony Mullen

Insight engines provide a range of capabilities to extract and contextualize data in support of digital dexterity and automation. Applications and software engineering leaders should view insight engines as a key component of digital workplace and hyperautomation.

Overview

Key Findings

- The range of embedded and custom touchpoints for the delivery of insight is increasing, serving multiple roles and use cases.
- Insight engines continue to vary significantly in terms of the range of prebuilt connectors, touchpoints and integrations available.
- Expansion to domain and situational use cases requires customization, which can be lengthy and challenging in the absence of productized solutions aligned to industry and functional verticals.

Recommendations

As an applications and software engineering leader responsible for search and insight in the context of the digital workplace, you should:

- Place the beneficiaries — people and machines — at the center of your selection and deployment initiative by defining success in terms of the business capabilities to be enabled and ways of working that will support these.
- Serve digital dexterity of users and process automation by working with stakeholders to identify the data sources, touchpoints and integrations required.
- Clarify what is out of the box versus what you or others must develop by treating insight engines as a platform upon which search and insight applications are provided or developed.

Strategic Planning Assumption(s)

By 2023, 60% of organizations will seek composability in new application investments.

By 2023, 85% of AI solutions by vendors will focus on concrete domains and industry verticals.

By 2025, 80% of large enterprises will need to have a conversational-technology-focused center of excellence or skills resource.

What You Need to Know

This report is aimed at applications and software engineering leaders responsible for enabling search and insight as part of their digital workplace (serving employees), digital experiences (serving customers) and digital platforms (serving machines). It will help them gain a deeper understanding of the products/services offered by the vendors in the insight engines market. It will also assist them in differentiating products as part of software selection or review of a current solution, or to support horizon scanning.

The comparative analysis provided here scores competing products/services against a set of critical differentiators – “critical capabilities” – identified by Gartner for this market’s principal use cases, namely:

- Internal (intranet) search
- External (website) search
- Extraction of data for analytics and automation
- Insight applications

We have attributed a weight to each capability for each use case to represent buyers in general. Use the “customize” feature on the interactive version to decide your own weightings, or add additional use cases. In addition to the vendors compared in this report, we encourage you to consider others that may provide solutions suited to your industry or specific requirements, such as those listed in the “Honorable Mentions” section of [Magic Quadrant for Insight Engines](#).


Analysis

Critical Capabilities Use-Case Graphics

Vendors' Product Scores for Internal (Intranet) Search Use Case



Product or Service Scores for Internal (Intranet) Search




Gartner

Source: Gartner (March 2021)

Vendors' Product Scores for External (Website) Search Use Case



Product or Service Scores for External (Website) Search




Source: Gartner (March 2021)

Vendors’ Product Scores for Extraction of Data for Analytics and Automation Use Case



Product or Service Scores for Extraction of Data for Analytics and Automation




Gartner

Source: Gartner (March 2021)

Vendors' Product Scores for Insight Applications Use Case



Product or Service Scores for Insight Applications



Gartner

Source: Gartner (March 2021)

Vendors

Coveo

Coveo's product is Coveo Relevance Platform, with variants based on the principal in-app experiences it supports, namely: Coveo for Salesforce, Coveo for ServiceNow, Coveo for Sitecore, Coveo for Salesforce B2B Commerce, and Coveo for AppDirect.

Coveo focuses on the productization of prebuilt integrations with principal applications, where domain- and situation-specific work takes place — for example, Salesforce in the context of customer support. The Coveo Relevance Platform is best suited to internal (intranet) search, with its ranking in this use case elevated by its capability to deliver to various touchpoints. This is demonstrated through the product's use within third-party applications by a broad range of functional roles across industries.

The platform can also be applied to the external (website) search, extraction of data for analytics and automation, and insight applications use cases. The product's lowest-scoring capabilities were query input flexibility and data enrichment.

Elastic

Elastic's product is Elastic Enterprise Search. This product suite is composed of three products: Workplace Search, App Search (for building search capabilities within applications) and Site Search (for external [website] search).

Elastic is unique in this group of vendors in providing an insight engine — Elasticsearch — that forms the core of multiple products (including IntraFind and Squirro) and optionally enables Coveo's prebuilt capability.

Elastic Enterprise Search is best suited to internal (intranet) search, with its ranking for this use case elevated by its capabilities to evaluate and tune relevance, and deliver to various touchpoints. This is demonstrated through the product's support for third-party applications via ready-made integrations such as Atlassian Confluence.

The product can also be applied to external (website) search, extraction of data for analytics and automation, and insight applications use cases. Its lowest-scoring capabilities were query input flexibility and data enrichment.

EPAM

EPAM's product is InfoNgen.

The product supports traditional enterprise search and knowledge discovery, as well as more tailored offerings for brand, competitive and market intelligence. A major differentiator for InfoNgen is its ability to use not just internal enterprise sources, but also enrich search experiences and development with external content — 200,000 web/social sources, 30,000 government and regulatory sources, as well as over 35 prebuilt industry taxonomies.

InfoNgen is best suited to extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets. This is demonstrated by the product's application to automation use cases involving, for example, robotic process automation (RPA), where prebuilt integrations with RPA market leaders are provided by the vendor.

InfoNgen can also be applied to the internal (intranet) search and insight applications use cases. The product was not evaluated for its application to the external (website) search. Its lowest-scoring capabilities were personalization and data enrichment.

Expert.ai

Expert.ai's product is NL Suite. It includes capabilities from what was previously Cogito Intelligence Platform and comprises five products: Discover, Studio, Answers, and industry-specific products for Insurance and Life Sciences.

The platform offers a broad set of natural language processing (NLP) capabilities, founded on semantic enablement through its integral knowledge graph. A key differentiator is the accelerated customization of the knowledge graph, over which it provides extensive control.

NL Suite is best suited to extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets. This is demonstrated through the product's use in the analysis of externally sourced data for decision support when applied to use cases such as competitive intelligence, price monitoring and social media analysis.

The product can also be applied to the extraction of internal (intranet) search and insight applications use cases. It was not evaluated for its application to the external (website) search. The product's lowest-scoring capabilities were personalization, delivery to various touchpoints and query input flexibility.

Funnelback

Funnelback's product is also named Funnelback.

Its principal capabilities span site search, intranet search, case management, knowledge management and product management (retail). Funnelback has packaged solutions for higher education, finance and membership organizations with strong features tailored to government (extract, transform, load [ETL], accessibility auditing, content life cycle). Funnelback's technology is bundled as an integrated component of the Squiz digital experience platform (DXP) or offered separately as Funnelback, a wholly owned subsidiary of Squiz.

The product is best suited to external (website) search, with its ranking in this use case elevated by its personalization capability, and its ability to evaluate and tune relevance. This is demonstrated through the product's customer-facing concierge services and embedding of content recommendations within search results.

Funnelback can also be applied to the internal (intranet) search, extraction of data for analytics and automation, and insight applications use cases. The product's lowest-scoring capabilities were data enrichment and ease of use for administrators and SMEs.

Google

Google's product is Google Cloud Search.

The product is provided as part of Google Workspace or can be acquired and run separately. This flexibility enables it to be applied as a foundational insight engine serving all employees, whether integrated with Google Workspace or used separately. In addition, its customizability enables application to domain and situational scenarios, serving specific cohorts. In this way, it provides a strong default for Google Workplace customers and a viable alternative for subscribers to other cloud office suites.

Google Cloud Search is best suited to insight applications, with its ranking in this use case elevated by its personalization capability, and ability to include key data sources. This is

demonstrated through the product's ability to support customized interfaces via its cloud search widget or the cloud search query API.

The product can also be applied to the internal (intranet) search and extraction of data for analytics and automation use cases. Google Cloud Search was not evaluated for its application to the external (website) search. The product's lowest-scoring capabilities were architecture and deployment model and query input flexibility.

IBM

IBM's product is IBM Watson Discovery. This supersedes the vendor's Watson Explorer product, which was based on acquired technology and deployed on-premises.

Watson Discovery has a strong focus on supporting customer care and customer support use cases. It deeply integrates with related offerings such as IBM Graph and IBM Watson Assistant. This enables the product's smart document understanding capabilities to be applied across a more comprehensive set of tools to deliver insight to employees.

The product is best suited to internal (intranet) search, with its ranking in this use case elevated by its capabilities to deliver to various touchpoints, and to evaluate and tune relevance. This is demonstrated through the product's ability to be integrated with other applications in order to deliver information in the context of the tool(s) used to progress work.

IBM Watson Discovery can also be applied to the extraction of data from analytics and automation and insight applications use cases. The product was not evaluated for its application to external (website) search. Its lowest-scoring capabilities were ability to include key data sources, and architecture and deployment model.

IHS Markit

IHS Markit's product is Goldfire.

Goldfire leverages rich NLP capabilities for structured documents to extract rich metadata managed in a separate semantic index. Through extensive faceting, this enables precise querying of large and complex corpi, which is well-suited to technical professionals.

Goldfire is best suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to include key data sources. This is demonstrated by its ability to support technical staff such as engineers and researchers working with complex documentation and datasets.

Goldfire can also be applied to the internal (intranet) search and insight applications use cases. The product was not evaluated for its application to external (website) search. Its lowest-scoring capabilities were architecture and deployment model, personalization, query input flexibility and security.

IntraFind

IntraFind's product is iFinder. The product is one of four in this Critical Capabilities report to be built on an open-source search engine, which is Elasticsearch in IntraFind's case.

iFinder differentiates with strong integrations into third-party semantic and graph tools such as Neo4j, TigerGraph and PoolParty to extend its core capabilities. It also rates among the highest in this group of evaluated solutions for flexible deployment methods of all or part of the solution.

The product is best suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets. This is demonstrated through the product's application to use cases such as product data search (PDM/PLM), threat analysis and web monitoring. IntraFind's iFinder product can also be applied to the internal (intranet) search and insight applications use cases. The product was not evaluated for its application to the external (website) search, although it does offer a site search product. Its lowest-scoring capabilities were query input flexibility and delivery to various touchpoints.

Lucidworks

The analysis below reflects Lucidworks' new capabilities scores since the initial publication of this research on 30 March 2021.

Lucidworks' product is Lucidworks Fusion. It is one of four products in this Critical Capabilities research to be built on an open-source search engine, which in Lucidworks' case, is Apache Solr. In addition, the vendor offers add-on applications: Smart Answers (conversational middleware for chatbots and virtual assistants) and Predictive Merchandiser (a business-facing tool that optimizes digital commerce experiences).

Fusion enables organizations to rapidly develop their own applications, beyond interfaces, for the delivery of insight. This helps reach domain and situational use cases, beyond a foundational offering for search and insight.

Fusion is best-suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets and data enrichment capability. This is demonstrated through the product's processing of signal and other data sources to support automation in the context of the digital commerce and digital workplace use cases.

Fusion can also be applied to the internal (intranet) search, external (website) search and insight applications use cases. The product's lowest-scoring capabilities were security and architecture and deployment model.

Since the initial publication of this Critical Capabilities document, Lucidworks has introduced Experience Optimizer in September 2021. This new feature elevates Fusion's score for the evaluate and tune relevance capability. Experience Optimizer enables line-of-business workers to evaluate and tune relevance in pursuit of optimal results, affecting internal (intranet) search and external (website) search use cases in particular.

Micro Focus

Micro Focus' product is IDOL, an acronym for Intelligent Data Operating Layer.

IDOL differentiates primarily in the space of government and military applications of insight, where well-developed security and compliance capabilities make it a strong choice. Also, its virtual assistant extension can support conversational question answering in natural language from explicit FAQ sources, “fact banks” or extracted from documents.

IDOL is best suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets and data enrichment capability. This is demonstrated through the product’s assimilation into other Micro Focus products to provide analytical capabilities, and application to a broad range of analytical use cases including rich-media analytics, social media analytics and scientific research.

IDOL can also be applied to the internal (intranet) search and insight applications use cases. The product was not evaluated for its application to the external (website) search. Its lowest-scoring capabilities were architecture and deployment model, and ease of use for administrators and SMEs.

Microsoft

Microsoft’s product is Microsoft Search, which is integral to Microsoft 365 and built on Microsoft Graph.

Given the prominence of Microsoft 365, the vendor’s approach is to integrate its offering into this product suite, rather than offering prebuilt — or the option of building, via APIs — integrations with other platforms and applications. This positions Microsoft Search for foundational application by all employees rather than domain or situational application by specific cohorts to specific contexts.

The product is best suited to insight applications, with its ranking in this use case elevated by its personalization capability, and ability to include key data sources. Unlike other products in this market, this is specific to Microsoft 365, and demonstrated through the integration of the product into a range of Microsoft 365 products.

Microsoft Search can also be applied to the internal (intranet) search use case. The product was not evaluated for its application to the external (website) search and extraction of data for analytics and automation use cases. Its lowest-scoring capabilities were architecture and deployment model, and delivery to various touchpoints.

Mindbreeze

Mindbreeze’s product is Mindbreeze InSpire.

InSpire offers the longest list of connectors to data sources in this Critical Capabilities report, enabling a broad range of content and data to be indexed.

The product is best suited to insight applications, with its ranking in this use case elevated by its capabilities in data enrichment and multiple languages. This is demonstrated through its application to a broad range of use cases and the customer-facing marketplace for insight applications developed by other customers.

Mindbreeze InSpire can also be applied to the internal (intranet) search and extraction of data for analytics and automation use cases. It received the highest score out of all vendor products in the three use cases for which it was evaluated. The product's lowest-scoring capabilities were query input flexibility, ease of use for administrators and SMEs, and delivery to various touchpoints.

Sinequa

Sinequa's product is Sinequa ES.

Sinequa builds on its strength in the analysis of content and data to synthesize, as well as surface, information to deliver insight. This, in turn, builds on an entity-centric perspective (which the vendor and others in the market term as 360 search).

The product is best suited to insight applications, with its ranking in this use case elevated by its multiple languages and personalization capabilities. This is demonstrated through its use across industries and functions to build custom insight applications serving a broad range of use cases.

Sinequa ES can also be applied to the internal (intranet) search, external (website) search, and extraction of data for analytics and automation use cases. The product's lowest-scoring capabilities were ease of use for administrators and SMEs, and data enrichment.

Squirro

Squirro's product is Squirro Insights Engine.

Beyond offering a platform on which insight applications can be developed for and by clients, Squirro develops and productizes insight applications for specific industries and functional domains. This enables clients to deliver directly to specific domain or situational use cases for which the vendor offers products. At present, this aligns to a range of specific industries and/or functional domains.

Squirro Insights Engine is best suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets. This is demonstrated through its focus on actionable, role-centric dashboards derived from various sources of content and data, including external news feeds.

The product can also be applied to the internal (intranet) search and insight applications use cases. It was not evaluated for its application to the external (website) search use case. The product's lowest-scoring capabilities were evaluate and tune relevance, and personalization.

Context

Although most organizations provide search-based insight, its usefulness is limited to a subset of data sources, excludes rich-media content, is unlikely to be embedded across various enterprise UIs, and lacks customization to the knowledge domain of end users.

Insight engines offer a single product to leverage the value of content and data to support multiple use cases across digital workplace, digital commerce and hyperautomation. Whether one product supporting all use cases, or an additional product complementing a foundational product to serve

domain and situational scenarios, insight engines offer maximum value when one or few are deployed. Key to this are multiple connectors to data sources, multiple touchpoints and integrations to people and machines, and a common index around which all capabilities are anchored.

As an applications and software engineering leader responsible for search and insight technologies, you should use this document only after you have identified your desired business capabilities, the ways of working that will enable these, and the people as well as machines that will deliver this work. This research will help you identify how insight engine products and their capabilities align with your organization's goals and priorities. You will then be able to build a shortlist of vendors that can help you achieve your objectives.

To simplify selection, and to help ensure that you make the right choice, consider your organization's longer-term objectives and the pace at which it wants to achieve them. This will help you narrow the selection and home in on those vendors that can provide the technology to support the associated solution(s). To narrow your selection further, focus not only on what is extremely important to your organization, but also on what is differentiating. The critical capabilities featured in this research reflect such areas, while the use cases represent scenarios in which such areas of primary and secondary differentiation play a critical role. Finally, combine the current research with the companion [Magic Quadrant for Insight Engines](#) to understand key market dynamics, together with how the featured vendors are scored and positioned in the context of those dynamics.

Product/Service Class Definition

In broad terms, an insight engine will possess the following characteristics:

- A platform composed of proprietary and (optionally) third-party components on which applications, or components of third-party applications, are provided or developed.
- Connectors to many and varied data sources located both on-premises and in the cloud, containing many and varied types of content and data.
- One or more indexes that serve as a structured data source registering documents, records and other digital assets with associated metadata.
- A service-oriented architecture that enables the integration of third-party components to enhance the capabilities of the platform (e.g., speech to text) or deliver components to other applications.
- The capability to develop integrations via APIs and web services and (optionally) packaged integrations to notable third-party applications (e.g., robotic process automation).
- A configurable pipeline for the preprocessing of content to extract data for indexing.
- A configurable query processor that accepts queries from users and machines to identify and return relevant entries from the index(es).

- Configurable interfaces through which queries are captured and results returned, enabling the delivery of content, data and information via many and varied touchpoints.
- The use of AI to enhance the accuracy and tolerance of the index, pipeline and query processor.

Critical Capabilities Definition

Ability to Analyze Result Sets

The ability to produce results in the form of data derived from indexed content, and to analyze this either in place or via syndication.

Once extracted from content, structured data can be analyzed either within the insight engine or via a separate tool. This enables insight engines to serve as a means to prepare data for analysis and utilize appropriate tools or processes according to use case. Products in this market vary according to the tools provided within the product, as well as the tools they can syndicate to outside the product.

Ability to Include Key Data Sources

The breadth of data sources and types that can be connected to and indexed comprehensively.

Content resides in a wide range of repositories and other data sources. These can be internal to the organization (e.g., content services platforms) or external (e.g., social media feeds), and contain many and varied content types (e.g., videos). Including key data sources requires “connectors” that allow data sources to be accessed and interrogated so their data can be transferred, processed and indexed for query. Products vary in this market according to the range of connectors available and how they are provided (by the vendor or third parties; included, free or priced), and according to the content types that can be indexed and how this is supported.

Architecture and Deployment Model

A model for architecture and deployment that works for organizations, including hosting, integration with other systems and extensibility.

Many factors affect where and how an insight engine is deployed, including the location of data sources to be indexed, governance and the touchpoints used by users. As platforms, the components of each product may be proprietary, third party or openly sourced, and integrated with other platforms, applications and components elsewhere in the wider digital ecosystem. A choice of on-premises or cloud-based deployment gives clients choice, while hybrid deployments further enhance adaptation to a client’s particular circumstances. Products in the market vary according to the ways in which the component parts are sourced, deployed and integrated.

Data Enrichment

The ability to model data, with or without the addition of external data, to create a richer index against which queries can be run.

Beyond being treated as a bag of words, content can be processed to extract data contained within it (e.g., reference to entities that are salient to the business or augmented to include data that is missing, such as classification using a controlled vocabulary). Data can also be remodeled. This enables a richer index that not only points to content that is relevant to queries (e.g., a document in a shared drive), but also becomes a source of content in and of itself (e.g., a knowledge graph supporting expertise location). Products in this market vary according to the range of techniques available (e.g., sentiment analysis, extraction of text from speech, entity extraction).

Delivery to Various Touchpoints

Flexibility to control the display of results and innovate new forms of presentation, and within applications other than the insight engine (e.g., CRM, custom-made applications).

Insight engines deliver insight by informing both surfacing content and data in context. Insight engines offer the capability to control context through both the query and the display of results, which may be an out-of-the-box application, a custom application or embedded within another application, explicitly or implicitly. Products in this market vary according to the provision of out-of-the box applications/UIs, APIs and web services, software development kits (SDKs), prebuilt integrations and application/UI development tools.

Ease of Use (Admins and SMEs)

The ability to enable workers who must administer the solution to ensure it is managed, developed and configured.

Control of insight engines via back-end interfaces enables their ongoing delivery and the efficacy of the service they provide. Administrators are enabled by the user experience (UX) of the UI. Extending this to a wider community of users – principally subject matter experts (SMEs) – enables the breadth of control necessary. This is especially so with respect to the ways in which relevancy is assessed and adjusted. Products in the market vary according to the level of control the UI provides, the level of UX and the roles involved.

Evaluate and Tune Relevance

The ability to assess the precision and recall of results and adjust configuration to optimize.

Insight engines use relevancy methods to match queries from users and machines against an index. Unlike databases, which structure and maintain data for precise recall, insight engines rely on relevancy methods to deliver accurate recall by creating structure where this is possible; compensating for its absence where not. Evaluating the performance of these methods and adjusting them to improve performance enables organizations to address the principal complaint of irrelevant results arising from imprecise and incomplete results. Products in the market vary according to level of automation, control and other techniques.

Multiple Languages

The ability to parse and process text from multiple languages, as well as provide UX (UI and documentation) in multiple languages.

Insight engines need to parse, process and index content encoded in multiple languages using multiple character sets. The value that can be obtained from indexed content depends on the depth of linguistic analysis available: morphology (words), syntax (sentences and clauses), semantic (meaning) and pragmatic (conversation). Insight engines must also support employees working in multiple languages, who need to use back-end UIs, read documentation and provide touchpoints to diverse user bases. Products in this market vary according to which languages are supported, the depth of linguistic analysis for each language, and whether the UI and documentation can be expressed in each language.

Personalization

The ability to vary results based on parameters captured from users and their touchpoints to personalize experiences.

Tailoring the query submitted and the results returned according to the user and their context enables more relevant results. Parameters including usage history, touchpoint, point in digital journey, geography, device and role enable users to be categorized so as to adjust their experience. This capability includes proactive delivery of information via recommendations, smart alerts and notifications. Products in this market vary according to the range of parameters they can utilize, the use of usage data, the scope for variation of queries and results, and other aspects.

Security

The ability to enable administrators to control access and authorization and comply with standards, regulations and laws.

Comprehensive support is needed to securely connect to multiple systems for the ingress, storage and egress of content and data in ways that honor a rich tapestry of authentication and authorization across content, repository, user, group and touchpoint. In addition to the standards/technologies for authentication and authorization (e.g., SAML), support for standards (e.g., SOC 2) and law (e.g., HIPAA) applies. Products in this market vary according to the standards/technologies and regulations/laws supported.

Query Input Flexibility

The ability to enable users to express their query in various ways and guide them through the querying process to deliver relevant responses.

Flexibility in the query process and method of expression facilitates the use of insight engines across domains, technologies and users; for example, natural language questions, images, whole documents, or aiding traditional queries with autocomplete. This is key to effective application to initial and future use cases, where contextualization is both core and essential. Query input flexibility also includes the ability of queries from machines (e.g., via API). Products in this market vary according to the capture of implicit (e.g., location) and explicit (e.g., keyword) parameters, the format these can take (e.g., audio or text) and proactive guidance (e.g., autosuggest/complete functionality).

Use Cases

Internal (Intranet) Search

Enterprisewide search within the digital workplace.

The principal use case for most organizations, and the default “application” for the insight engine products provided by vendors, is enabling all employees to find files and records located within many and various repositories, databases and other data sources. Taken collectively, these data sources may contain thousands, millions or even billions of individual files and records.

External (Website) Search

Site search within and across the enterprise’s public-facing and multichannel digital presence.

Search within websites — and related channels — is powered by products in the insight engines and digital commerce search markets, with some overlap between the two. Site search is principally concerned with indexing content served via HTTP(S) and with an emphasis on customized presentation to facilitate seamless UX. More advanced usage enables interaction with the content and functionality of digital commerce, web content management and digital experience platforms.

Extraction of Data for Analytics and Automation

Deriving data from indexed content for the purposes of analysis and reporting, either within the insight engine or via syndication to other applications.

With connectors to a wide range of data sources, insight engines can index the full spectrum of content and data. The ability to extract data (tokens modeled for machines) from content (tokens modeled for humans) enables this data to be analyzed, both directly and in combination with other data. In this way, data is prepared for analysis. This “data analysis” can take place within the product — using tools that are provided or which can be added — or outside the product via syndication to other tools. As such, insight engines can serve as a solution or a component of a larger solution for analysis, or delivery of data to machines to serve automation.

Insight Applications

Delivery of insight to digital workers within the context of custom-made or third-party applications.

When considered as a platform, insight engines provide the opportunity to develop multiple custom-made applications. Such applications, developed and deployed on top of the core insight engine platform, can serve a variety of use cases in the context of domains defined by industry vertical and business function. Although typically developed by or for clients, some vendors have productized such insight applications. Examples include course finders (for universities), expertise directories and knowledge bases.

Vendors Added and Dropped

Added

- Elastic

■ Squirro

Dropped

- **Attivio:** This vendor’s IP, core technology and a selection of key staff have been acquired by ServiceNow. Attivio continues to serve customers it acquired previously, but no longer competes meaningfully in this market.
- **Dassault Systèmes:** A design and engineering specialist offering a wide array of product development services across manufacturing and 3D products and environments. Although Dassault Systèmes offers a comprehensive set of solutions for industrial applications, the applicability of its insight engine is limited for other purposes and beyond its other products.

Inclusion Criteria

To qualify for inclusion in this Critical Capabilities report, vendors needed to:

- Sell a product that:
 - Met the market definition: Insight engines are systems that apply relevancy methods to describe, discover, organize, and analyze content and data. This approach allows for existing or synthesized information to be delivered proactively or interactively in the context of digital workers, customers or other constituents at timely business moments.
 - Was available as a separate, independent software product or as part of a cloud office platform intended to support knowledge workers in all roles.
 - Used connectors to access and gather content/data of varying types from multiple repositories other than those sold by the vendor.
 - Created an index by processing gathered content/data, and made this available via query through multiple touchpoints (for humans) and integrations (for machines).
 - Served two of the four principal use cases for products in this market, namely internal (intranet) search and insight applications.
 - Had been generally available for at least a year, as of midnight, U.S. Eastern Daylight Time on 1 April 2020.
 - Had technical capabilities, features and functionality present within the product or supported for download prior to midnight, U.S. Eastern Daylight Time on 1 April 2020.

Table 1: Weighting for Critical Capabilities in Use Cases

--	--

Critical Capabilities ↓	Internal (Intranet) Search ↓	External (Website) Search ↓	Extraction of Data for Analytics and Automation ↓	Insight Applications ↓
Ability to Analyze Result Sets	9%	9%	12%	9%
Ability to Include Key Data Sources	13%	10%	13%	13%
Architecture and Deployment Model	6%	6%	6%	6%
Data Enrichment	9%	9%	12%	12%
Delivery to Various Touchpoints	12%	6%	9%	9%
Ease of Use (Admins and SMEs)	6%	6%	6%	6%
Evaluate and Tune Relevance	12%	12%	9%	6%
Multiple Languages	9%	12%	9%	12%
Personalization	9%	12%	9%	12%

Critical Capabilities ↓	Internal (Intranet) Search ↓	External (Website) Search ↓	Extraction of Data for Analytics and Automation ↓	Insight Applications ↓
Query Input Flexibility	9%	9%	9%	9%
Security	6%	9%	6%	6%
As of February 2021				

Source: Gartner (March 2021)

This methodology requires analysts to identify the critical capabilities for a class of products/services. Each capability is then weighted in terms of its relative importance for specific product/service use cases.

Each of the products/services that meet our inclusion criteria has been evaluated on the critical capabilities on a scale from 1 to 5.

Critical Capabilities Rating

Table 2: Product/Service Rating on Critical Capabilities

Critical Capabilities ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelb
Ability to Analyze Result Sets	2.9	3.6	3.6	3.2	2.9
Ability to Include Key Data Sources	2.7	2.8	1.6	1.7	2.8
Architecture and Deployment Model	2.8	2.4	1.8	2.0	2.6

Critical Capabilities ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelb
Data Enrichment	2.0	2.2	1.3	1.8	1.3
Delivery to Various Touchpoints	3.8	3.0	1.4	1.2	2.8
Ease of Use (Admins and SMEs)	3.6	2.9	1.6	2.7	2.0
Evaluate and Tune Relevance	3.2	3.0	3.3	1.8	3.9
Multiple Languages	2.9	3.6	3.5	2.4	2.8
Personalization	3.8	3.0	1.0	1.0	4.0
Query Input Flexibility	2.0	1.3	1.9	1.2	2.5
Security	3.2	3.9	2.8	2.3	3.2
As of February 2021					

Source: Gartner (March 2021)

Table 3 shows the product/service scores for each use case. The scores, which are generated by multiplying the use-case weightings by the product/service ratings, summarize how well the critical capabilities are met for each use case.

Table 3: Product Score in Use Cases

Use Cases ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelback
Internal (Intranet) Search	2.99	2.87	2.16	1.87	2.85
External (Website) Search	2.98	2.92	N/A	N/A	2.90
Extraction of Data for Analytics and Automation	2.93	2.86	2.17	1.93	2.78
Insight Applications	2.95	2.86	2.10	1.88	2.78
As of February 2021					

Source: Gartner (March 2021)

To determine an overall score for each product/service in the use cases, multiply the ratings in Table 2 by the weightings shown in Table 1.

Critical Capabilities Methodology

This methodology requires analysts to identify the critical capabilities for a class of products or services. Each capability is then weighted in terms of its relative importance for specific product or service use cases. Next, products/services are rated in terms of how well they achieve each of the critical capabilities. A score that summarizes how well they meet the critical capabilities for each use case is then calculated for each product/service.

"Critical capabilities" are attributes that differentiate products/services in a class in terms of their quality and performance. Gartner recommends that users consider the set of critical capabilities as some of the most important criteria for acquisition decisions.

In defining the product/service category for evaluation, the analyst first identifies the leading uses for the products/services in this market. What needs are end-users looking to fulfill, when

considering products/services in this market? Use cases should match common client deployment scenarios. These distinct client scenarios define the Use Cases.

The analyst then identifies the critical capabilities. These capabilities are generalized groups of features commonly required by this class of products/services. Each capability is assigned a level of importance in fulfilling that particular need; some sets of features are more important than others, depending on the use case being evaluated.

Each vendor's product or service is evaluated in terms of how well it delivers each capability, on a five-point scale. These ratings are displayed side-by-side for all vendors, allowing easy comparisons between the different sets of features.

Ratings and summary scores range from 1.0 to 5.0:

1 = Poor or Absent: most or all defined requirements for a capability are not achieved

2 = Fair: some requirements are not achieved

3 = Good: meets requirements

4 = Excellent: meets or exceeds some requirements

5 = Outstanding: significantly exceeds requirements

To determine an overall score for each product in the use cases, the product ratings are multiplied by the weightings to come up with the product score in use cases.

The critical capabilities Gartner has selected do not represent all capabilities for any product; therefore, may not represent those most important for a specific use situation or business objective. Clients should use a critical capabilities analysis as one of several sources of input about a product before making a product/service decision.

History

Vendors

The following providers have updated capacity scoring since the initial publication of this research 30 March 2021. This section features the original write-up and capability scores for those providers that have been updated.

Lucidworks

Lucidworks' product is Lucidworks Fusion. It is one of four products in this Critical Capabilities report to be built on an open-source search engine, which is Apache Solr in Lucidworks' case. In addition, the vendor offers add-on applications: Smart Answers (conversational middleware for chatbots and virtual assistants) and Predictive Merchandiser (a business-facing tool to optimize digital commerce experiences).

Fusion enables organizations to rapidly develop their own applications, beyond interfaces, for the delivery of insight. This helps reach domain and situational use cases, beyond a foundational offering for search and insight.

Fusion is best suited to the extraction of data for analytics and automation, with its ranking in this use case elevated by its ability to analyze result sets and data enrichment capability. This is demonstrated through the product's processing of signal and other data sources to support automation within the context of digital commerce and digital workplace use cases.

Fusion can also be applied to the internal (intranet) search, external (website) search and insight applications use cases. The product's lowest-scoring capabilities were security and architecture and deployment model.

Critical Capabilities Rating as of 9 February 2021

Table 4 and 5 featured in the original iteration of this research, which was published on 30 March 2021. Each of the products/services were evaluated on the critical capabilities on a scale of 1 to 5; a score of 1 = Poor (most or all defined requirements are not achieved), while 5 = Outstanding (significantly exceeds requirements). Table 5 shows the product/service scores for each use case. They featured in the original iteration of this research, which was published on 30 March 2021. The scores, which are generated by multiplying the use-case weightings by the product/service ratings, summarized how well the critical capabilities were met for each use case.

Table 4: Product/Service Rating on Critical Capabilities

Critical Capabilities ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelb
Ability to Analyze Result Sets	2.9	3.6	3.6	3.2	2.9
Ability to Include Key Data Sources	2.7	2.8	1.6	1.7	2.8
Architecture and Deployment Model	2.8	2.4	1.8	2.0	2.6
Data Enrichment	2.0	2.2	1.3	1.8	1.3

Critical Capabilities ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelb
Delivery to Various Touchpoints	3.8	3.0	1.4	1.2	2.8
Ease of Use (Admins and SMEs)	3.6	2.9	1.6	2.7	2.0
Evaluate and Tune Relevance	3.2	3.0	3.3	1.8	3.9
Multiple Languages	2.9	3.6	3.5	2.4	2.8
Personalization	3.8	3.0	1.0	1.0	4.0
Query Input Flexibility	2.0	1.3	1.9	1.2	2.5
Security	3.2	3.9	2.8	2.3	3.2
As of February 2021					

Source: Gartner (March 2021)

Table 5: Product Score in Use Cases

Use Cases ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelback

Use Cases ↓	Coveo ↓	Elastic ↓	EPAM ↓	Expert.ai ↓	Funnelback
Internal (Intranet) Search	2.99	2.87	2.16	1.87	2.85
External (Website) Search	2.98	2.92	-1.00	-1.00	2.90
Extraction of Data for Analytics and Automation	2.93	2.86	2.17	1.93	2.78
Insight Applications	2.95	2.86	2.10	1.88	2.78
As of February 2021					

Source: Gartner (March 2021)

**Learn how Gartner
can help you succeed**

Become a Client

© 2022 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's research organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources

believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner research may address legal and financial issues, Gartner does not provide legal or investment advice and its research should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its research is produced independently by its research organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)."

[About](#) [Careers](#) [Newsroom](#) [Policies](#) [Site Index](#) [IT Glossary](#) [Gartner Blog Network](#) [Contact](#) [Send Feedback](#)



© 2022 Gartner, Inc. and/or its Affiliates. All Rights Reserved.